

Remarks:

Pending in this application at the time of the February 24, 2005 Office Action were claims 1-57. By this amendment, Applicant has (1) canceled claims 3-4, 12-13, 26-27, 33-34, 39-40, and 47-48, (2) amended claims 2, 5-11, 14-18, 20, 23, 25, 29, 32, 42, and 54-55, and (3) added new claims 58-88. Thus, the currently pending claims are claims 1-2, 5-11, 14-25, 28-32, 35-38, 41-46, and 49-88. No new matter is present herein. Support for the new and amended claims can be found throughout the application. (See, e.g., Application at page 13, line 27- page 14, line 6; page 16, line 9 - page 21, line 35; Figures 4(a), 5, 10, and 11).

The outstanding Office Action rejected claims 1- 46 and 49-55 for obviousness based on a variety of combinations of the Goldberg patent (US Patent No. 6,389,038), the Brech patent (U.S. Patent No. 5,754,768), the Carlson patent (US Patent No. 6,298,070), the Lee patent (U.S. Patent No. 5,191,578) and the Arnold patent (U.S. Patent No. 5,199,028). The Office Action further rejected claims 47 and 48 under 35 U.S.C. 102(e) as being anticipated by the Goldberg patent, and allowed claims 56 and 57. Applicant submits claims 56 and 57 in their original form. By this amendment, Applicant has canceled claims 47 and 48 without prejudice, thereby rendering the anticipation rejection based on Goldberg moot. Applicant respectfully disagrees with the aforementioned obviousness rejections for the reasons stated below.

Applicant first notes that the claimed invention specifically relates to **a packet switch** and the **internal operations** of a packet switch that combines data packets into train packets according to the data packets' switch processing parameters (SPPs). As explained in the application at page 1, line 1 through page 4, line 24, this invention is addressed to solving the problem of scaling packet switches to handle higher rates of packet traffic. As line rates and/or the number of input ports to a packet switch have increased, a problem has developed wherein switch fabrics have less and less time to perform arbitration between the various incoming data packets that are simultaneously competing for access to the switch fabric outputs. Arbitration (referred to in the specification as "global processing") involves, inter alia, (1) processing the SPPs of the incoming packets that are simultaneously competing for access to an output side of the switch fabric to determine which packet will be selected for the next routing through the switch fabric and determine the switch fabric output port to which the selected packet will be routed, (2) configuring a path within the switch fabric to route the selected packet to the

appropriate switch fabric output port, and (3) actually routing the selected packet through the configured path.

The higher the line rate of packets handled by a conventional packet switch, the less time that the conventional packet switch will have to perform arbitration. For example, if the packet switch requires more time to perform arbitration than the time interval between receipt of consecutive packets, then the packet switch will become a traffic bottleneck. This can be a particularly acute problem for packet switches that employ bit slicing (see Figures 3(a) and (b) of the application), because the slicing operation further reduces the time interval between consecutive packets that must be processed through the switch fabric.

As a solution to this scaling problem, the inventors herein have developed a packet switch that queues packets according to their SPPs, creates train packets from the queued packets, and routes the train packets through the switch fabric. Because the train packets will often have a length that is greater than the individual data packets from which the train packets are created, the packet switch will have more time to perform arbitration (global processing) because the time interval between consecutive train packets can be expected to be longer than the time interval between consecutive individual data packets. Furthermore, because of the minimum length that can be defined for the train packets, the packet switch can be provided with a guaranteed minimum time interval for performing arbitration. Taken together, the present invention thereby allows packet switch designers to scale packet switches to higher speeds while still preserving the ability to implement desired arbitration schemes. Without the valuable contribution of the present invention, it is believed by the inventors herein that packet switch designers will often be forced to "strip down" the arbitration techniques implemented by the packet switch to a possibly undesired quality level.

The references cited in the Office Action, when considered alone or in combination, fail to disclose, teach, or suggest such an invention. Applicant will now address the shortcomings of these references relative to the claimed invention.

I. The cited references fail to disclose, teach or suggest a packet switch that queues train packets according to their SPPs to create train packets therefrom and then routes the train packets through a switch fabric according to the train packet's SPP.

The Office Action rejected independent claims 1, 19, 38, 49, and 52-54 for obviousness based on the combination of the Brech and Goldberg. Brech discloses a system whereby at the receiving end of a network, individual packets are combined based on the their session identifiers (session IDs) to form train packets. A train packet comprising a plurality of packets sharing a common session ID is then provided to a host processor as a group. The goal of the Brech system is to reduce the number of interrupts for the host processor as it receives packets. (See Brech; col. 1, 67 – col 2, line 9). Thus, with the Brech system, instead of the host processor experiencing 4 interrupts as 4 data packets are provided to it (even if those 4 packets share a common session identifier), the Brech system would reduce the number of interrupts to 1 (by way of example) by combining those 4 packets into a single packet train. Accordingly, the Brech system implements packet training in a completely different environment (within a receiving station 100 rather than a packet switch) on the basis of different parameters (a session ID rather than an SPP) and for different purposes than the claimed invention (to reduce the number of host processor interrupts rather than providing a packet switch with more time to perform arbitration).

Claim 1 recites the step of “routing each train packet through the switch fabric as specified by its encapsulated SPP”. Claim 19 recites the limitation “wherein the switch fabric is configured to route each received train packet to a switch fabric output according to the SPP included in the header of each train packet”. Claim 38 recites “a controller configured to (1) queue each data packet in an appropriate waiting buffer according to its SPP such that data packets sharing a common SPP are commonly-queued, and (2) create train packets from the commonly-queued data packets, each train packet having a payload and a header, wherein the payload of at least one train packet is comprised of a plurality of commonly-queued data packets, and wherein the header of each train packet includes the SPP corresponding to each data packet within the payload of that train packet”. Similar limitations are found in independent claims 49 and 52-54 as well as new claim 76. The Office Action contends that the “session ID” used by Brech is an equivalent to the switch processing parameter (SPP) used

by the claimed invention. (See, e.g., Office Action, page 4, lines 2-3). Applicant respectfully disagrees. First, ***the train packets created in the Brech system are not routed through a switch fabric at all***, and they are particularly not routed through a switch fabric on the basis of their session IDs. For any packet switching that may occur within the various networks disclosed by Brech, ***this packet switching occurs on individual packets and not on the packet trains***. Instead, Brech teaches that packet training should occur at a receiving station within the network. With the Brech system, packets are grouped at a receiving station according to their session IDs to create packet trains. These packet trains are in turn provided to a host processor 116 via a system bus 118 (see Brech; Figure 2; see also step 222 in Figure 6). Accordingly, the Brech patent fails to disclose, teach, or suggest to a person having ordinary skill in the art that a packet's session ID should be used as a switch processing parameter for a train packet.

Further still, a person having ordinary skill in the art would not be motivated by the Brech patent to incorporate SPP-based train packet queuing into a packet switch. The goal of packet training in the Brech system is to reduce the number of times that a host processor is interrupted. However, this problem does not correlate to the environment of a packet switch where the concern is increasing the amount of time available for arbitration. Thus, a person having ordinary skill in the art would not be motivated by the Brech patent to find a solution to the problem of providing packet switches with more time for arbitration.

The Goldberg reference, when considered alone or in combination with Brech reference, fails to bridge the gaps left by Brech. Goldberg teaches that SuperPackets should be created from individual packets to thereby reduce the ratio of header information to payload data in point-to-point network communications, particularly node-to-node Voice-over-Internet packet (VOIP) communications. (See Goldberg; col. 1, line 66 – col. 2, line 5; “[i]t is an object of the present invention to address the ***inefficient use of bandwidth resulting from communication of protocol headers*** in the transmission of packets over a common communications channel. It is a further object of the present invention to provide ***a more efficient UDP/IP header structure*** for transmitting packets over a common IP link” (emphases added); see also col. 1, lines 39-52; col. 3, lines 10-40). The Goldberg reference is silent with respect to performing

packet switching for packets arriving on multiple switch fabric inputs and competing for access to the output-side of the switch fabric.

Figures 6(a) and (b) of Goldberg illustrate the purpose behind the Goldberg system. In Figure 6(a), 4 individual data packets are shown, each within its own header H and payload V. As explained in an example at col. 3, lines 10-40, the header information H for each packet may account for 28 bytes of data and the payload V for each packet may account for 24 bytes. If these 4 packets are transmitted over a network individually, the ratio of payload-to-header (representing a measure of bandwidth efficiency) would be approximately 46%. However, Goldberg teaches that packets with sufficiently similar headers can be combined to form a SuperPacket, *wherein the individual headers of the individual packets are discarded*, the packet payloads are concatenated within a SuperPacket payload, and wherein *a new SuperPacket header (SPH) is attached to the SuperPacket*. This new SuperPacket header is described by Goldberg as *requiring fewer bytes* than the aggregated number of bytes in the 4 individual data packet headers H1 through H4, *thereby increasing the payload-to-header ratio for the packets communicated over a network*. It is for this purpose that the Office Action contends it would be obvious to combine Goldberg with Brech to arrive at the claimed invention. (See Office Action, page 4; "[t]he motivation would be to decrease the amount of space devoted to the header information in the superpacket.").

However, with the claimed invention, the SPP-based train packet queuing will often actually *decrease* the ratio of header information to payload information in the packets that are routed across the switch fabric. An example will help illustrate this point. Assume a train packet has payload length of 128 bytes. Further assume the packet switch receives a 1280 byte long packet. Further assume that the train packet header is 5 bytes. With the SPP-based train packet queuing of the present invention, the 1280 byte packet will be broken into 10 train packets. Thus, when switching the 10 train packets through the switch fabric, the total number of bytes processed by the switch fabric would be $128 \times 10 + 5 \times 10 = 1330$ bytes. However, if this 1280 byte packet was processed by the switch fabric without the SPP-based packet training of the present invention, the total number of bytes processed by the switch fabric would only be $1280 + 5 = 1285$ bytes. *Thus, the SPP-based train packet queuing of the*

present invention actually would operate to reduce the bandwidth efficiency of the packet switch, in direct contradiction to the express teaching of the Goldberg patent.

Accordingly, because the SPP-based train packet queuing of the present invention would result in a packet switch that operates expressly against the teachings of the Goldberg reference, a person having ordinary skill in the art ***would not be motivated*** by Goldberg to arrive at the claimed invention ***because doing so would have required that ordinarily-skilled person to proceed contrary to Goldberg's teachings.*** (See In re Gordon, 733 F. 2d 900, 902 (Fed. Cir. 1984) (holding that it is inappropriate to reject a claim for obviousness where the invention disclosed in the cited reference would be rendered inoperable for its intended purpose if used in the claimed invention). Rather than operating to increase the bandwidth efficiency of the packet switch, the SPP-based train packet queuing of the present invention instead operates to increase the amount of time available to the packet switch to perform global processing.

Therefore, based on these shortcomings, the combination of the Brech and Goldberg references fails to render concept of using SPP-based train packet queuing within a packet switch obvious. Brech is silent on the subject of group packets by their switch processing parameters (SPPs) and then routing the grouped packets through the switch fabric as a train packet according to the grouped packet's common SPP. Instead, Brech relies upon a different aspect of a packet to perform grouping and fails to teach that packet switching should be performed on the grouped packet trains. Moreover, the motivation behind the Goldberg SuperPacket is to reduce the header-to-payload ratio of packets sent across a network. With the claimed invention however, it is the case that this ratio may be increased rather than decreased. Proceeding contrary to the teachings of Goldberg, the claimed invention instead performs SPP-based train packet queuing to increase the amount of time generally available to the packet switch to perform arbitration (and in instances where a minimum train packet length is specified, define a minimum amount of time for the packet switch to perform arbitration). As such, Applicant respectfully submits that the obviousness rejection of independent claims 1, 19, 38, 49, and 52-54 should be withdrawn. Applicant further submits that new independent claim 76 is patentable over the Brech/Goldberg combination for the same reasons.

II. Claims 10, 28, 51, and 55 are also not rendered obvious by the cited combination of the Brech, Goldberg, and Lee references.

The Office Action rejected independent claims 10, 28, 51, and 55 for obviousness based on the Brech/Goldberg/Lee combination. As explained above in Section I, the Brech/Goldberg combination fails to disclose the SPP-based train packet processing technique of the present invention. It thus follows that this combination fails to disclose, teach or suggest the variations of this technique disclosed in the application (i.e., sequential train packet processing, sequential-to-parallel train packet processing, cocktail train packet processing, and hierarchical train packet processing) as each of these specific embodiments requires SPP-based train packet queuing. The Office Action cites the Lee patent for the proposition of processing sliced train packets through a plurality of parallel switch planes. However, the Lee reference is silent on the subject of train packet processing and therefore fails to bridge the shortcomings relative to the claimed invention left by Brech and Goldberg. As such, Applicant respectfully submits that the obviousness rejection of claims 10, 28, 51, and 55 is improper and should be withdrawn. Moreover, Applicant notes that the Office Action's citation to the Goldberg reference as teaching train packet slicing for the packet switch of the present invention is inappropriate. Goldberg teaches that "[w]hen the number of channels to be combined causes the total number of bytes to exceed the maximum wide area network packet size, the SuperPackets are divided into smaller SuperPackets which are below the maximum packet size." (See Goldberg; col. 7, lines 47-51). However, this teaching is inapplicable to the packet switch environment of the present invention. ***Within the packet switch***, the network's maximum packet size is ***not*** a concern. The train packet size can readily exceed the network's maximum packet size because, on the output side of the switch fabric, a train packet that exceeds the network's maximum packet size will be broken back down into its constituent data packets before those constituent data packets are sent out across the network. Therefore, a person having ordinary skill in the art would not be motivated by this teaching in the Goldberg reference to perform slicing upon the train packets of the present invention prior to processing by a stacked switch fabric.

III. Claims 58, 61, 64-66, 70-72, 87, 88 are patentable over the cited references because the cited references fail to disclose, teach or suggest the use of SPP mappers or a classification step to determine SPPs for data packets to processed through a packet switch that employs SPP-based train packet processing.

The Office Action cites the Brech reference for teaching the use of SPPs for train packets. However, as explained above in Section I, the Brech reference in fact fails to teach or suggest the use of SPPs for train packets. Accordingly, Applicant respectfully submits that claims 58, 61, 64-66, 70-72, 87, 88 are patentable over the cited references because these claims also include limitations relating to the SPP mapping operations of the present invention.

IV. Claims 59, 62, 68, 74, 82, and 83 are patentable over the cited references because the cited references fail to disclose, teach or suggest the use of a timer threshold configured to control how much padding will be processed through a switch fabric within train packets.

Claims 59, 62, 68, 74, 82 and 83 recite the use of a timer threshold to effectively control how much padding will be processed through a switch fabric within the train packets. The patent application describes this timer threshold as Timer1 in Figure 10 and the corresponding text in the specification. The cited references are silent on such a feature of train packet processing. Accordingly, Applicant respectfully submits that these claims are patentable over the cited references.

V. The cited references do not disclose, teach or suggest traffic distributors that sort and process incoming data packets into distribution classes on the basis of a predetermined criteria or a method of sorting incoming data based on such distribution classes.

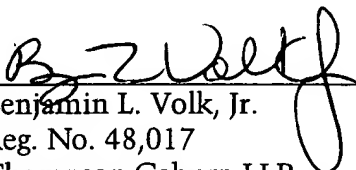
With respect to claims 8, 9, 17, 18, 49, 50, 51, 52 and 53, the cited references fail to disclose, teach or suggest the limitations relating to traffic distributors that sort and process data packets into distribution classes on the basis a predetermined criteria in combination with the performance of SPP-based train packet processing on the sorted data packets (or a method of sorting incoming packets based such distribution classes in combination with the performance of SPP-based train packet processing on the sorted data packets). The Brech reference fails to disclose the aforementioned limitations relating to the combination of SPP-based train packet processing with traffic distributors and the method of sorting data packets

according to distribution classes. As noted in Section I, Brech discloses a method of grouping packets based on the session ID between two hosts or logical units; however, there is no disclosure in Brech relating to a method of sorting and processing data based on distribution classes. Brech also fails to disclose traffic distributors capable of sorting and processing packets into such distribution classes. Contrary to the assertion in the Office Action, a person having ordinary skill in the art would not be motivated to find the traffic distributors of the claimed limitations above. As such, Applicant respectfully submits that claims 8, 9, 17, 18, 49, 52 and 53 are not rendered obvious by the cited references.

Conclusion:

For the foregoing reasons, Applicant respectfully submits that the currently pending claims are allowable over the cited prior art. Applicant respectfully submits that other dependent claims of the patent application not specifically addressed herein are also independently patentable over the cited references. However, because arguments herein establish the nonobviousness of one or more antecedent claims for these dependent claims, this response does not specifically address these dependent claims' patentability. Moreover, Applicant in no way concedes that the references cited by the Office Action constitute prior art and reserves the right to later establish, if necessary and appropriate, that one or more of the cited references are not in fact prior art under U.S. patent law. Favorable action is respectfully requested.

Respectfully submitted,


Benjamin L. Volk, Jr.
Reg. No. 48,017
Thompson Coburn LLP
One US Bank Plaza
St. Louis, Missouri 63101
(314) 552-6000
(314) 552-7000 (fax)